

Solving the TTC 2011 Reengineering Case with GrGen.NET

Edgar Jakumeit Sebastian Buchwald

Karlsruhe Institute of Technology (KIT)

buchwald@kit.edu

The challenge of the Reengineering Case [4] is to extract a state machine model out of the abstract syntax graph of a Java program. The extracted state machine offers a reduced view on the full program graph and thus helps to understand the program regarding the question of interest. We tackle this task employing the general purpose graph rewrite system GrGen.NET (www.grgen.net).

1 What is GrGen.NET?

GRGEN.NET is an application domain neutral graph rewrite system [1], the feature highlights regarding practical relevance are:

Fully Featured Meta Model: GRGEN.NET uses attributed and typed multigraphs with multiple inheritance on node/edge types. Attributes may be typed with one of several basic types, user defined enums, or generic set, map, and array types.

Expressive Rules, Fast Execution: The expressive and easy to learn rule specification language allows straightforward formulation of even complex problems, with an optimized implementation yielding high execution speed at modest memory consumption.

Programmed Rule Application: GRGEN.NET supports a high-level rule application control language, Graph Rewrite Sequences (GRS), offering sequential, logical, iterative and recursive control plus variables and storages for the communication of processing locations between rules.

Graphical Debugging: GRShell, GRGEN.NET's command line shell, offers interactive execution of rules, visualizing together with yComp the current graph and the rewrite process. This way you can see what the graph looks like at a given step of a complex transformation and develop the next step accordingly. Or you can debug your rules and sequences.

2 The Core Assignment

The task of the core assignment is to extract a state machine model out of the abstract syntax graph of a Java program. The task stems from the domain of reengineering where software engineers need to gain insights into legacy systems, which is a lot easier given a birds eye view on the high level structure (and thus behavior) of the program. The aim of the task is to evaluate the solutions and the tools backing them regarding performance and scalability, with a domain leading naturally to large graphs to be considered; and especially to evaluate the solutions and tools regarding the ability to – and conciseness in – carrying out complex, non-local matchings of graph elements, requiring the matching of recursive graph structures. Before the extraction can take place, the Java program graph needs to be imported from an Ecore file describing the source metamodel and an XMI file specifying the graph. Afterwards the resulting state machine has to be exported into an XMI file conforming to a given Ecore file describing the state machine metamodel.

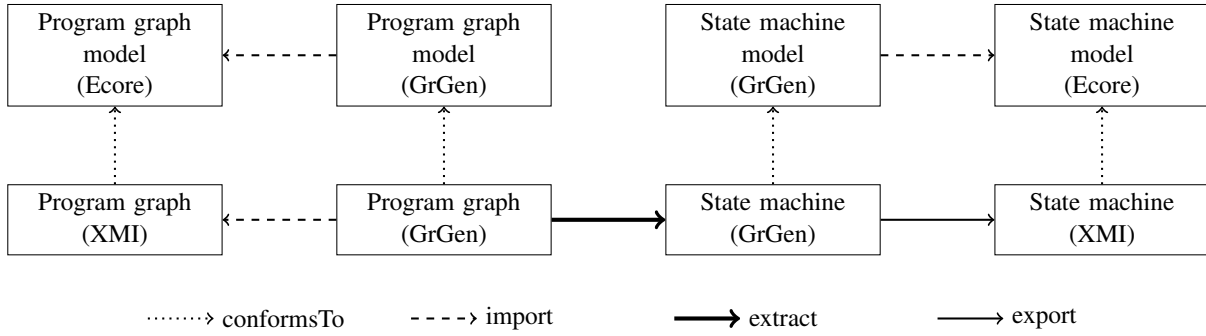


Figure 1: Processing steps of the model extraction. The extraction and the XMI export are written in GrGen.NET languages. Import is handled by a supplied import filter, which generates GrGen meta model files (.gm) as an intermediate step.

2.1 Importing the Graph

As GrGen.NET is a general purpose graph rewrite system and not a model transformation tool, we do not support importing Ecore metamodels directly. Instead, we supply an import filter generating an equivalent GrGen-specific graph model (.gm file) from a given Ecore file by mapping classes to GrGen node classes, their attributes to corresponding GrGen attributes, and their references to GrGen edge classes. Inheritance is transferred one-to-one, and enumerations are mapped to GrGen enums. Class names are prefixed by the names of the packages they are contained in to prevent name clashes; the same holds for references which are prefixed by their node class name. Afterwards the instance graph XMI adhering to the metamodel described in the Ecore file thus adhering to the just generated equivalent GrGen graph model is imported by the filter into the system to serve as the host graph for the following extractions, i.e. transformations. The entire process is shown in Figure 1 above.

2.2 Extracting the States

The transformation is done in two steps, the first creates the states of the state machine, and the second inserts the transitions in between. Each step consists of the application of one rule (utilizing a subpattern) on all matches found, giving a direct correspondence between coding conventions and rules.

Let us start with a short introduction into the syntax of the basic constructs of the rule language: Rules in GrGen consist of a pattern part specifying the graph pattern to match and a nested rewrite part specifying the changes to be made. The pattern part is built up of node and edge declarations or references with an intuitive syntax: Nodes are declared by $n:t$, where n is an optional node identifier, and t its type. An edge e with source x and target y is declared by $x -e:t-> y$, whereas $-->$ introduces an anonymous edge of type `Edge`. Nodes and edges are referenced outside their declaration by n and $-e->$, respectively. Attribute conditions can be given within `if`-clauses.

The rewrite part is specified by a `replace` or `modify` block nested within the rule. With `replace`-mode, graph elements which are referenced within the `replace`-block are kept, graph elements declared in the `replace`-block are created, and graph elements declared in the pattern, not referenced in the `replace`-part are deleted. With `modify`-mode, all graph elements are kept, unless they are specified to be deleted within a `delete()`-statement. Attribute recalculations can be given within an `eval`-statement. These and the language elements we introduce later on are described in more detail in our solution of the Hello World case [2], and especially in the extensive GrGen.NET user manual [1].

Now let us have a look at the code to create the states (here and in the following rules the prefixes from name mangling were removed due to space constraints):

```
rule createStates
{
  stateClass:Class;
  stateClass -:annotationsAndModifiers-> :Abstract;
  if { stateClass.name == "State"; }

  es:CreateStates(stateClass);

  modify {
    sm:StateMachine;
    es(sm);
  }
}
```

We search for the abstract class of name `State` as starting point and create the state machine which will receive the states and transitions found. The real work is done in a subpattern `CreateStates`, of which an instance `es` is declared and thus searched from the found `stateClass` on; or better in the rewrite part of this subpattern, which is applied with rule call syntax passing the just created `StateMachine` node:

```
pattern CreateStates(parentClass:Class) modify(sm:StateMachine)
{
  iterated {
    extendingClass:Class -:extends-> r:NamespaceClassifierReference;
    r -:classifierReferences-> cr:ClassifierReference;
    cr -:target-> parentClass;

    es:CreateStates(extendingClass);

    optional {
      negative {
        extendingClass -:annotationsAndModifiers-> :Abstract;
      }

      modify {
        sm -:states-> s:State;
        s -:link-> extendingClass;
        eval { s.name = extendingClass.name; }
      }
    }

    modify {
      es(sm);
    }
  }

  modify { }
}
```

The subpattern searches a class directly extending the given parent class passed as a parameter, i.e. it descends one subtyping step downwards in the type hierarchy. It matches into breadth with the `iterated` construct to get all such extending classes; the `iterated` matches all instances of its contained pattern which can be found in the host graph. Then the subpattern matches further into depth by call-

ing itself recursively with the just matched extending class as parameter, this way we cover the entire type hierarchy from the parent class on. In the optional case the class is not abstract a state is created within the state machine and a link edge is created linking the state with the class (the link edges are additional helper edges introduced to store the traceability relation between source and target nodes). The optional matches the contained pattern if it is available in the host graph. The negative causes matching of the containing pattern to fail if its pattern can be found in the host graph.

2.3 Extracting the Transitions

The transitions are inserted with a second rule given in Figure 2 utilizing the subpattern given in Figure 3. We search for the `class.Instance().activate()` pattern in the graph, if found we know the target state from the class of the called method and the link between the class and the state we inserted previously. Then we search with the subpattern `FindSourceState` for the source state, which gets yielded¹ into the `def` pattern element `sourceState`. If this was found we add a `Transition` in between the source state and the target state, and link it to the `expressionStatement` containing the method call.

The subpattern is used to recursively walk outwards from the method call to the class containing the call (recursive ascent over statement nesting); passing over the different types of statements and statement containers which might be on the way, until the class is reached yielding it back. The statements passed are all linked to the transition, which will be helpful for the extension tasks.

2.4 Extension Tasks

The trigger attributes of the transitions are filled by four rules for the four different ways specified; they get executed one after the other (this way handling the priority), first the non-run method, then the switch case, then the catch block and finally the fallback rule. Due to the links from the transitions to all the constructs on the path from the method call to the containing class this is a simple local pattern search, as can be seen in Figure 4.

The action attributes of the transitions are filled by two rules for the two different ways specified; first the enum value used in a send method, then the fallback rule. Again this is a simple local pattern search due to the links from the transitions to all the constructs on the path from the method call to the containing class, cf. Figure 5.

2.5 Exporting the State Machine

A visualization of the resulting state machine is given in Figure 6. The XMI export of this graph is handled by 5 additional rules given in `export.gri` containing `emit` statements: one for assigning XMI ids to the elements to be exported, which are stored in a map from the nodes to the corresponding ids, two for writing the XMI prefix and suffix, one for writing the `States` and one for writing the `Transitions`, utilizing the previously computed node to id mapping.

3 Rule Control, Performance, and Visualization

The extraction process is controlled by the graph rewrite script `reengineering.grs` which is executed by the GrShell; the script contains the following lines:

¹The `yield`-keyword is used to open a block with yielding assignments, and as a mandatory annotation at an entity whenever it is yielded to.

```

import primitive_types.ecore java.ecore StateMachine.ecore
1_small-model.xmi reengineering.grg

xgrs [createStates]
xgrs [createTransitions]
xgrs [addTriggerNonRunMethodName] ;> [addTriggerSwitchCaseEnumValueName] \
    ;> [addTriggerCatchBlockExceptionClassName] ;> [addTriggerOtherwise]
xgrs [addActionSend] ;> [addActionOtherwise]

```

The import command imports the XMI input graph complying to the Ecore models, and additionally includes the rules given in the rule file. The `xgrs` keyword starts an extended graph rewrite sequence, which is the rule application control language of GrGen (prepending `debug` before `xgrs` allows you to debug the sequence execution in GrShell). The rules are executed on all the matches found, which is requested by the all-bracketing `[rule]`. The then-right operator `>` executes the left sequences, then the right sequence, and returns as sequence result the result of the right sequence; the sequence results are irrelevant for this task, in general they are used to control sequence applications.

3.1 Performance

The benchmark results for the extraction task are given in the following table.

set no.	import time	import size	shell time	shell size	extraction time
1	2,855	2.0	31	3.5	130
2	2,917	2.1	32	3.6	140
3	17,878	188.8	4,165	420.9	187
1	1,279	1.3	46	2.3	125
2	1,314	1.3	47	2.4	130
3	28,658	105.3	7,800	277.4	213

Table 1: Results for different input sets; running time in ms, memory usage in MiBytes.

The reported values are computed as the arithmetic mean of the middle 3 values out of 5 measurements, on a Core i7 920 (2.6GHz) with 6 GiBytes of main memory under Windows Vista 64 Bit with MS .NET 64 Bit for the upper part of the table and on a Core 2 Duo U9600 (1.6GHz) with 3 GiBytes of main memory under Windows 7 32 Bit with MS .NET 32 Bit for the lower part. Import time is the time needed for importing the graph, import size is the size of the heap after importing the graph. Shell time is the time needed to transform the imported graph from API level (as seen by an external program using the API) into a named graph as employed by the GrShell of the rapid prototyping environment, shell size is the size of the heap after the named graph was constructed. Extraction time is the time needed for the application of the extraction rules. Remark: the dominating component of the extraction time is the time needed by the .NET just-in-time compiler producing machine code out of the .NET bytecode.

3.2 Visualization

The GrShell utilizes the graph viewer yComp as its visualization component; the final state machine visualization was already presented in Figure 6. But in addition to the state machine, yComp is able to give a decent visualization of the original program graph, too, as you may see in Figures 7, 8 and 9 which give a series of images zoomed in, an outstanding help in understanding and debugging.

This is made possible by the high configurability of yComp: one can choose from several available layout algorithms, e.g. organic for a force-based layout, or hierarchic which works well for program graphs (it was used in rendering the syntax graphs shown in this paper). For every available node or edge type it can be configured in which color with what node shape or edge style it should be shown, with what attribute values or fixed text as element labels or tags it is to be displayed, or if it should be shown at all. Furthermore graph nesting can be configured by registering edges at certain nodes to define a containment hierarchy, causing the nodes to be displayed as subgraphs containing the elements to which they are linked by the given edges. In Figure 10 an excerpt from our configuration file for customizing the graph layout of the program graph is shown. In addition a helper edge introduction step was added, so that all expression nodes are nested inside their containing statements, not only the outermost ones. A helper step was used in producing the final state machine visualization Figure 6, too, replacing Transition nodes with real edges.

4 Conclusion

In this paper we presented a GrGen.NET solution to the Reengineering challenge of the Transformation Tool Contest 2011 (which is available as a SHARE image [5]). The abstract Java syntax graph conforming to the `java.ecore` metamodel was imported by a supplied import filter under remapping to the graph concepts supported by GrGen; the extracted state machine was exported by a handful of text emitting graph rewrite rules. The name mangling Ecore/XMI import and the explicitly programmed XMI export, which result from the fact that GrGen is not a dedicated model transformation tool, are the only points of our solution we regard to be rather weak. In contrast to the declarative graph rewriting with recursive and iterated structures (see [3] for more on this) which we employed to extract a state machine giving a high level overview out of the original syntax graph. This ability of matching and rewriting recursive patterns allowed us to give a concise and simple solution to the core task of the Reengineering challenge closely following the specification given, with rules matching kernel patterns and subpattern recursion and iteration to match recursive structures into depth and breadth. During rewriting of the recursive match for transition creation, from the activation call to the containing class outwards, links were inserted from the transition to the elements visited; besides having been a help in debugging they especially allowed to easily solve the extension task with purely local graph rewrite rules (we could have used storagemaps instead of storing and retrieving this traceability information as other tools have done, but using edges allows for easy, visual debugging). The extraction resulted in correct target state machines for all input graphs supplied; for the large graph it took about 200ms, i.e. performance is an order of magnitude better than the reference solution. The goal of the task is to allow program understanding by extracting and displaying a reduced, easily understandable model. In addition to visualizing this simple model we have presented a visualization of the original program graph with our graph viewer yComp; utilizing color customization and especially nesting subgraphs inside nodes, we managed to achieve an understandable visualization for the two medium sized example graphs (7000 nodes plus 7000 edges), which a reengineer can inspect and navigate in order to develop and debug his transformation rules.

References

- [1] Jakob Blomer, Rubino Geiß & Edgar Jakumeit (2011): *The GrGen.NET User Manual*. <http://www.grgen.net>.

- [2] Sebastian Buchwald & Edgar Jakumeit (2011): *Saying Hello World with GrGen.NET - A Solution to the TTC 2011 Instructive Case*. In Van Gorp et al. [6].
- [3] Berthold Hoffmann, Edgar Jakumeit & Rubino Geiß (2008): *Graph Rewrite Rules with Structural Recursion*. 2nd Intl. Workshop on Graph Computational Models (GCM 2008). "<http://www.informatik.uni-bremen.de/~hof/papers/08-GCM.pdf>".
- [4] Tassilo Horn (2011): *Program Understanding: A Reengineering Case for the Transformation Tool Contest*. In Van Gorp et al. [6].
- [5] Edgar Jakumeit & Sebastian Buchwald (2011): *SHARE demo related to the paper Solving the TTC 2011 Reengineering Case with GrGen.NET*. http://is.ieis.tue.nl/staff/pvgorp/share/?page=ConfigureNewSession&vdi=XP-TUe_TTC11_GrGen_v2.vdi.
- [6] Pieter Van Gorp, Steffen Mazanek & Louis Rose, editors (2011): *TTC 2011: Fifth Transformation Tool Contest, Zürich, Switzerland, June 29-30 2011, Post-Proceedings*. EPTCS.

A Code Listings and Screenshots

```

rule createTransitions
{
  expressionStatement:ExpressionStatement -:>expression-> refTargetClass;
  refTargetClass:IdentifierReference -:>target-> targetClass:Class;
  refTargetClass -:>next-> callInstance;
  callInstance:MethodCall -:>target-> instance:ClassMethod;
  callInstance -:>next-> callActivate;
  callActivate:MethodCall -:>target-> activate:ClassMethod;
  if { instance.name=="Instance" && activate.name=="activate"; }

  targetClass <-:link- targetState:State;
  def sourceState:State;
  fss:FindSourceState(expressionStatement, yield sourceState);

  sm:StateMachine;

  modify {
    sm -:>transitions-> transition:Transition;
    sourceState <-:src- transition -:>dst-> targetState;
    sourceState -:>out-> transition <-:in- targetState;
    transition -:>link-> expressionStatement;
    fss(transition);
  }
}

```

Figure 2: Rule for creating transitions

```

pattern FindSourceState(containedEntity:Node, def sourceState:State)
    modify(transition:Transition)
{
  alternative {
    StatementListContainer {
      listContainer:StatementListContainer -:statements-> containedEntity;
      fss:FindSourceState(listContainer, yield sourceState);
      modify {
        transition -:link-> listContainer;
        fss(transition);
      }
    }
    StatementContainer {
      container:StatementContainer -:statement-> containedEntity;
      fss:FindSourceState(container, yield sourceState);
      modify {
        transition -:link-> container;
        fss(transition);
      }
    }
    StatementSwitch {
      switch:Switch -:cases-> containedEntity;
      fss:FindSourceState(switch, yield sourceState);
      modify {
        transition -:link-> switch;
        fss(transition);
      }
    }
    StatementCondition {
      condition:Condition -:elseStatement-> containedEntity;
      fss:FindSourceState(condition, yield sourceState);
      modify {
        transition -:link-> condition;
        fss(transition);
      }
    }
    StatementTry {
      try:TryBlock -:catcheBlocks-> containedEntity;
      fss:FindSourceState(try, yield sourceState);
      modify {
        transition -:link-> try;
        fss(transition);
      }
    }
    Class {
      cc:Class -:members-> containedEntity;
      ss:State -:link-> cc;
      yield { yield sourceState = ss; }
      modify {
        transition -:link-> cc;
      }
    }
  }
}

```

Figure 3: Subpattern for creating transitions


```

rule addTriggerNonRunMethodName
{
  transition:Transition -:link-> method:ClassMethod;
  if { method.name != "run"; }

  modify {
    eval { transition.trigger = method.name; }
  }
}

rule addTriggerSwitchCaseEnumValueName
{
  transition:Transition -:link-> case:NormalSwitchCase;
  case -:condition-> caseCondition:IdentifierReference;
  caseCondition -:target-> value:EnumConstant;

  modify {
    eval { transition.trigger = value.name; }
  }
}

rule addTriggerCatchBlockExceptionClassName
{
  transition:Transition -:link-> catchBlock:CatchBlock;
  catchBlock -:parameter-> parameter:OrdinaryParameter;
  parameter -:typeReference-> nspClassRef:NamespaceClassifierReference;
  nspClassRef -:classifierReferences-> classRef:ClassifierReference;
  classRef -:target-> exceptionClass:Class;

  modify {
    eval { transition.trigger = exceptionClass.name; }
  }
}

rule addTriggerOtherwise
{
  transition:Transition;
  if { transition.trigger == null || transition.trigger == ""; }

  modify {
    eval { transition.trigger = "--"; }
  }
}

```

Figure 4: Rules for extension task 1

```
rule addActionSend
{
  transition:Transition -> link-> block:StatementListContainer;
  block -> statements-> exprStmt:ExpressionStatement;
  exprStmt -> expression-> callMethod:MethodCall;
  callMethod -> target-> method:ClassMethod;
  callMethod -> arguments-> enumClassRef:IdentifierReference;
  enumClassRef -> next-> enumValueRef:IdentifierReference;
  enumValueRef -> target-> enumValue:EnumConstant;
  if { method.name == "send"; }

  modify {
    eval { transition.action = enumValue.name; }
  }
}

rule addActionOtherwise
{
  transition:Transition;
  if { transition.action == null || transition.action == ""; }

  modify {
    eval { transition.action = "--"; }
  }
}
```

Figure 5: Rules for extension task 2

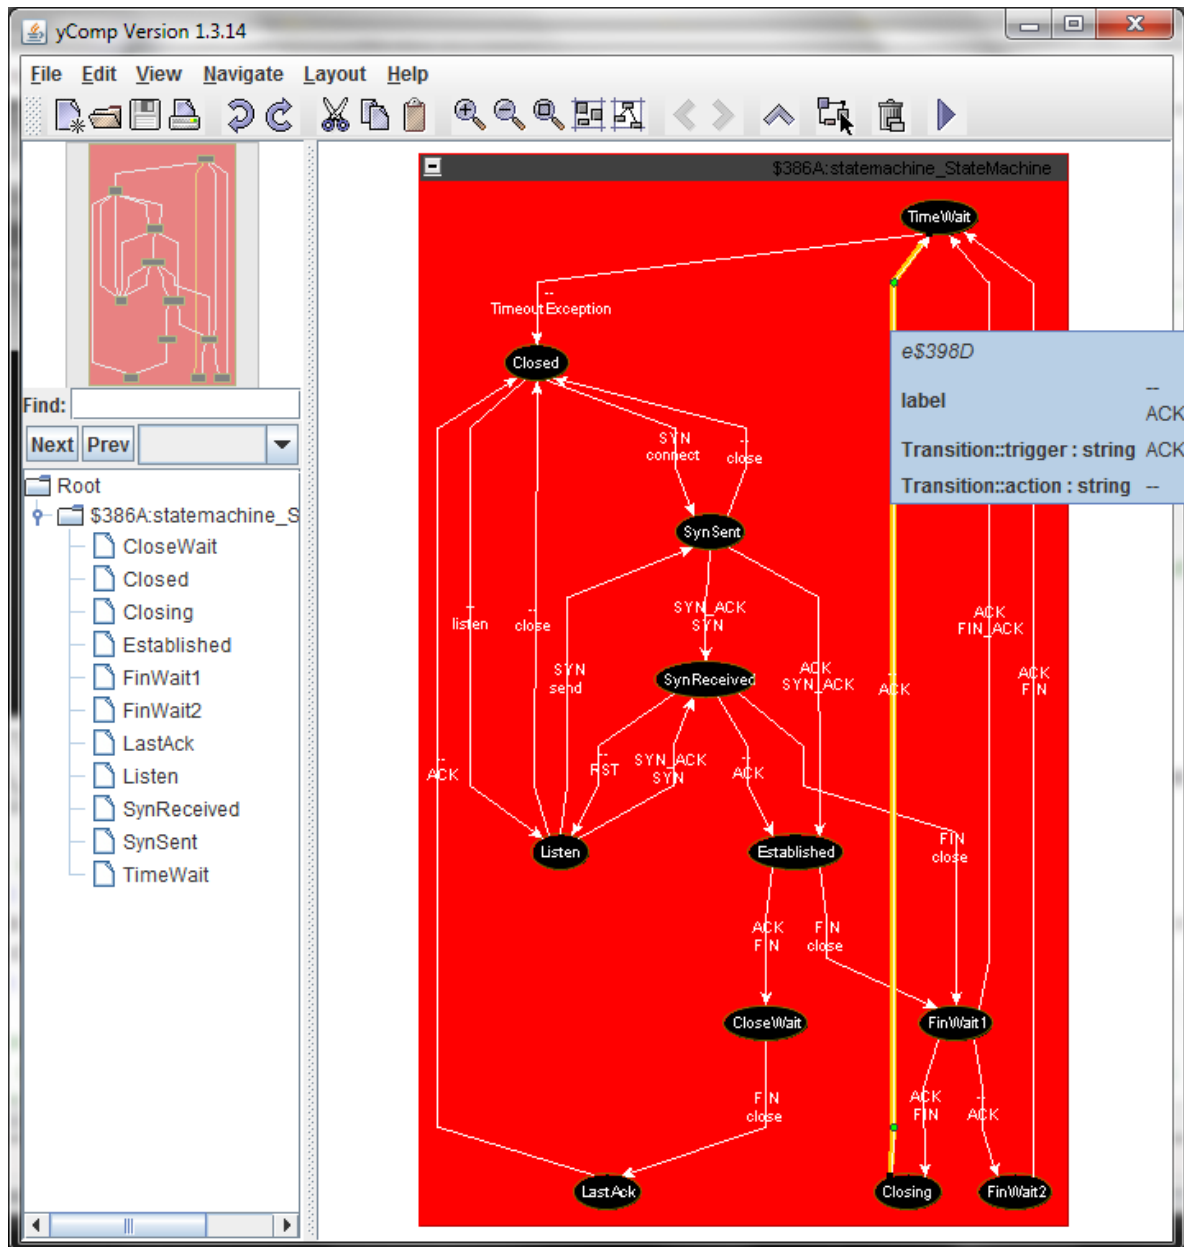


Figure 6: The resulting state machine, with an edge selected and its attributes displayed

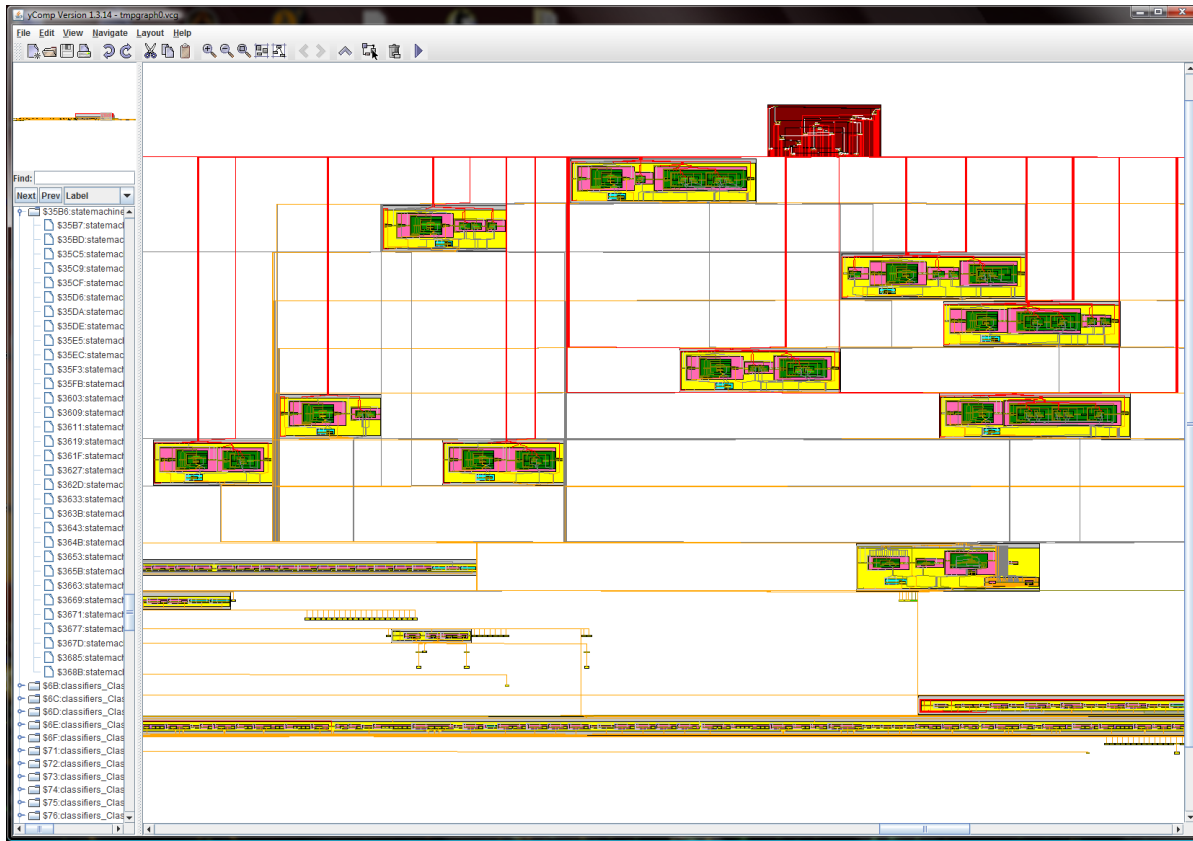


Figure 7: The program graph, the rectangles are the classes, the top one is the state machine

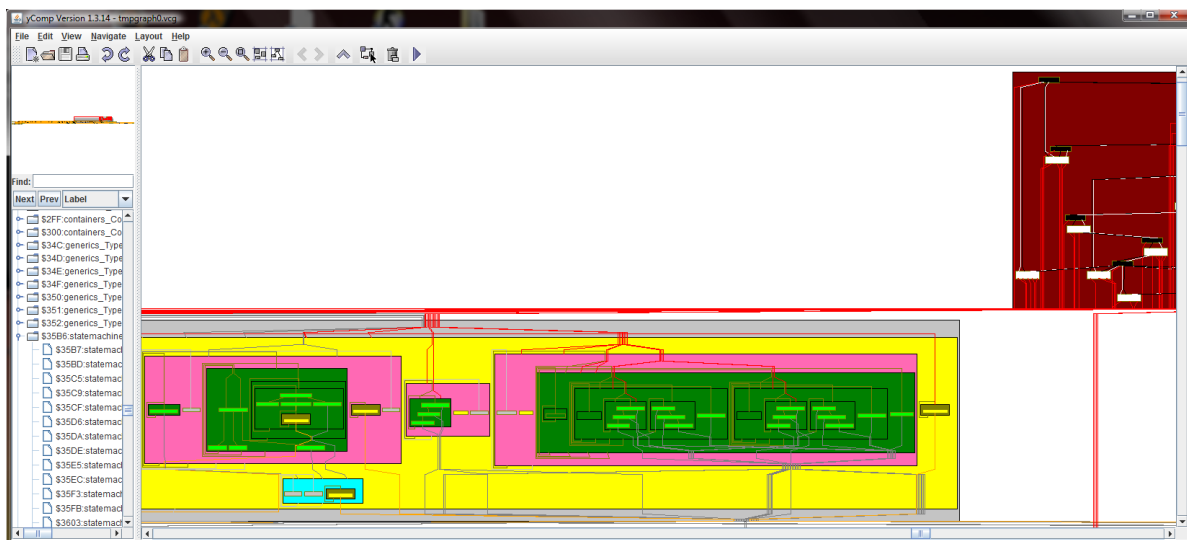


Figure 8: The program graph zoomed with the class SynSent and the state machine

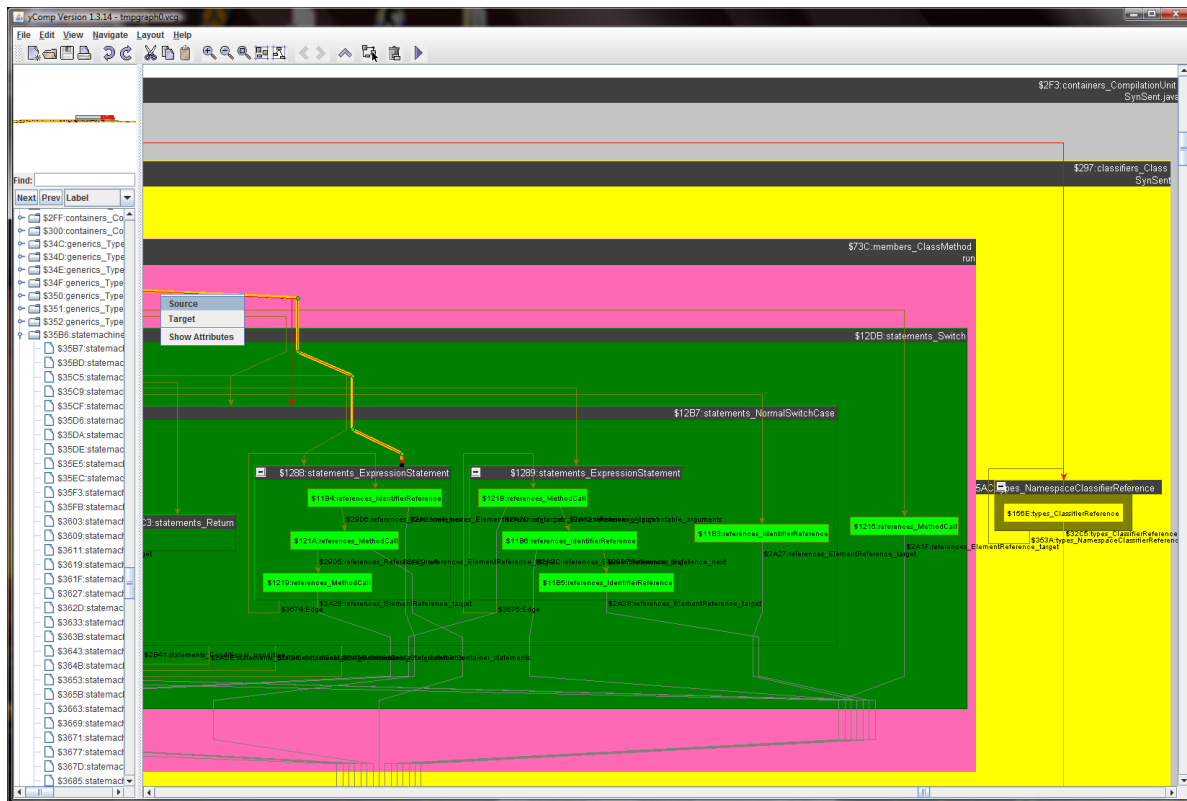


Figure 9: The program graph zoomed further to the method run of the class SynSent

```

debug set layout Hierarchic

dump add node classifiers_Classifier
    group by hidden outgoing members_MemberContainer_members
dump add node members_Method
    group by hidden outgoing members_MemberContainer_members

dump add node classifiers_Annotation exclude

dump set node members_Method color pink
dump set node members_Field color cyan

dump set edge references_ElementReference_target color grey

dump add node classifiers_Class shortinfotag _name
dump add node members_Method shortinfotag _name

```

Figure 10: An excerpt from the configuration of the layout